

Dimension Reduction for Physiological Variables Using Graphical Modeling

Dr. Michael Imhoff,

Surgical Department, Klinikum Dortmund, 44137 Dortmund, Germany

Dr. rer. nat. Roland Fried, Prof. Dr. rer. nat. Ursula Gather, Dipl.-Stat. Vivian Lanius,
Department of Statistics, University of Dortmund, 44221 Dortmund, Germany

Abstract

In intensive care, physiological variables of the critically ill are measured and recorded in short time intervals. The proper extraction and interpretation of the essential information contained in this flood of data can hardly be done by experience alone. Typically, decision making in intensive care is based on only a few selected variables. Alternatively, for a dimension reduction statistical latent variable techniques like principal component analysis or factor analysis can be applied. However, the interpretation of latent variables extracted by these methods may be difficult. A more refined analysis is needed to provide suitable bedside decision support. Graphical models based on partial correlations provide information on the relationships among physiological variables that is helpful for variable selection and for identifying interpretable latent components. In a comparative study we investigate how much of the variability of the observed multivariate physiological time series can be explained by variable selection, by standard principal component analysis and by extracting latent components from groups of variables identified in a graphical model.

Introduction

In intensive care, clinical information systems acquire and store physiological variables and device parameters online at least every minute. A physician may be confronted with more than 200 variables in the critically ill during typical morning rounds¹. Detection of critical states and of intervention effects based on these data is of great importance for bedside decision support. It should also be noted that constant information overload is one contributing cause for preventable medical errors^{2,3}.

A statistical analysis of the recorded data exhibits strong correlations e.g. between the variables of the hemodynamic system (different types of blood pressures, heart rate, pulse, and blood temperature) measured at short time intervals. Multivariate time series modeling should appropriately reflect the dependencies among these variables. This claim usually leads to complex models involving numerous parameters and requiring a high amount of data to enable reliable inference. Thus, suitable strategies for dimension reduction are also required when applying automatic

statistical techniques as the available data often does not suffice to model the full set of variables. This problem is known as the curse of dimensionality.

Thus, besides the aim of detecting changes in the patient's conditions, reducing the number of variables is a further task. Typically, some of the variables are selected according to personal experience. Of course, this is subjective, and it is important to know which and how much information we neglect in the reasoning process based on such a selection. The selection should be guided by information on the relationships between the variables. Statistical techniques like factor and principal component analysis allow to extract latent, i.e. unobservable variables which describe the correlations among the observed variables better and capture more of their variability than a simple variable selection. However, such latent variables are often not meaningful although it is important that they can be interpreted by healthcare professionals if we wish to make decisions about interventions or changes of treatments.

In order to overcome these difficulties we apply graphical interaction models. These models have become an important tool for investigating and modeling relationships in multivariate data as they allow a simple graphical visualization. The variables are represented by vertices and the relationships between the variables are illustrated by edges. Separations in the graph provide information about direct and indirect relationships in the data^{4,5,6,7}. In the following we exploit the separation properties of partial correlation graphs and relate them to dynamic factor models. This allows extraction of latent variables which are meaningful and explain more of the observed variability in the data than a simple variable selection.

Methods

Data set. On the surgical intensive care unit of the Klinikum Dortmund, a tertiary referral centre, online monitoring data was acquired from 25 consecutive critically ill patients (9 female, 16 male, mean age 66 years) with extended hemodynamic monitoring requiring pulmonary artery catheterization, in one minute intervals with a standard clinical information system. This data was transferred into a secondary SQL database and exported into standard statistical software for further analysis. A total of 129943 sets of observations were analyzed.

In the analysis we concentrated on the variables heart rate HR, pulse PULS, arterial diastolic pressure APD, arterial systolic pressure APS, arterial mean pressure APM, pulmonary artery diastolic pressure PAPD, pulmonary artery systolic pressure PAPS, pulmonary artery mean pressure PAMP, central venous pressure CVP and blood temperature Temp. In order to eliminate artifacts and irrelevant short-term fluctuations we removed outliers for each variable individually using a robust procedure based on the repeated median, which allows to preserve trends as well as sudden level shifts in the data^{8,9}. Thus we retain the relevant variability in the data when reducing the dimension, but not irrelevant outliers.

Partial correlation graphs. Between multiple variables usually a multitude of relationships exists, but many of them are indirect, i.e. they are induced by others. Distinguishing between direct and induced relationships among the observed variables is difficult from experience alone. Graphical models based on partial correlations reveal the essential relationships which are not induced by other variables. Visualization of a graphical model is accomplished by a graph: We draw a circle for each variable and connect each pair of variables by an undirected edge (a simple line) representing a symmetrical interaction whenever the relation between these variables persists after conditioning on all the other variables. Missing edges indicate the indirect character of some marginal relationships which are induced by underlying conditional dependencies. Indirect relationships can result from successively ordered direct influences. A subset of the variables is called complete if each pair of these variables is connected by an edge, i.e. if none of these relationships is indirect.

From a statistical point of view, measurements of physiological variables observed in short time intervals constitute multivariate time series as there may be interactions not only between instantaneous but also between time-lagged observations. Therefore we use partial correlation graphs for multivariate time series^{10,11,12}. Here, linear relationships between all pairs of variables at all time lags are investigated controlling for the linear effects of the other variables at all time lags, i.e. after all linear effects of the other series have been removed¹³. These relationships are called partial correlations and can be expressed equivalently in the frequency domain using the partial spectral coherency, that measures the partial correlations at all frequencies. Hence, partial correlation graphs allow to detect relationships in form of partial linear, possibly time-lagged dependencies between the variables of a multivariate time series. Moreover, under some weak regularity assumptions we can interpret the separations found in a graphical model. It has already been shown that "empirical relationships" found by partial correlation graphs correctly represent

"physiological relationships" based on medical knowledge within the hemodynamic system¹⁴.

Dynamic factor models and partial correlation graphs. Dynamic factor analysis allows to model a multivariate time series using a lower dimensional process of latent, i.e. unobserved variables called factors. A simple dynamic factor model for an observed multivariate time series $Y(t)$ is given by

$$Y(t) = \Lambda X(t) + \varepsilon(t)$$

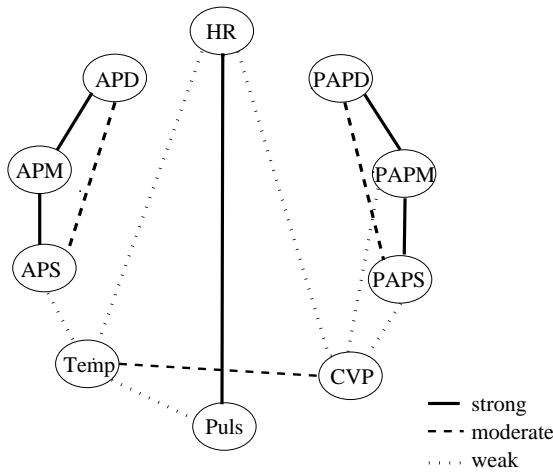
with an unobserved factor process $X(t)$ of latent variables and an error process $\varepsilon(t)$, each following a vector autoregressive model, which is a standard statistical model for multivariate time series data¹⁵. Here, Λ is a matrix of unknown parameters called loadings. If we observe d variables, i.e. $Y(t)$ is d -variate, and model the data using k latent variables, i.e. $X(t)$ is k -variate, $k < d$, then the loading at the (i,j) -th position in the $(d \times k)$ -matrix Λ describes the impact of the j -th factor on the i -th observed variable. If for a factor many of its loadings are close to zero we can identify it with the group of observed variables for which the loadings are large in absolute value. This model can be fitted to the observed data by analyzing the autocovariance matrices at the first few time lags similarly to performing an ordinary principal component analysis¹⁶. More general factor models can be formulated where $Y(t)$ is not only influenced by instantaneous factors $X(t)$, but also by time-lagged factors $X(t-h)$. However, identification and fitting such more general models is difficult since more parameters need to be estimated and since we need to decide which time lags should be considered.

Assuming that the spectral density matrix of the multivariate stationary time series $Y(t)$ is regular at all frequencies, an algorithm has been derived for construction of the partial correlation graph for the observable variables given an underlying factor model of very general form¹⁷. It turns out that in case of uncorrelated factors and uncorrelated error processes a pair of observed variables (A,B) is connected by an edge if and only if both variables have nonzero loadings for one of the factors, or a sequence of variables A_1, \dots, A_m exists such that all of the pairs (A, A_1) , (A_1, A_2) , ..., (A_m, B) fulfill the former condition. This allows to deduce suitable assumptions for a factor model from a preliminary data analysis using partial correlation graphs. Particularly, the resulting graph provides an assistance in identifying the number and types of factors. It seems straight forward to identify a complete subset in a partial correlation graph for the observable process with a latent factor. However, the identification of such common factors can be obscured since dependencies within the error process $\varepsilon(t)$ or between the factors can result in additional edges in the partial correlation graph. Nevertheless, it seems reasonable to attribute strong relations to the factors and weaker ones to the errors.

Results

In order to get a general impression about the relationships between the hemodynamic variables we constructed a partial correlation graph for each patient. The program Spectrum18 was used for the calculations. A typical example of such a graph resulting from a joint analysis of the partial spectral coherencies between all pairs of variables is shown in Figure 1. We use different edges for indicating different strengths of relationships as measured by the area below the partial spectral coherencies¹².

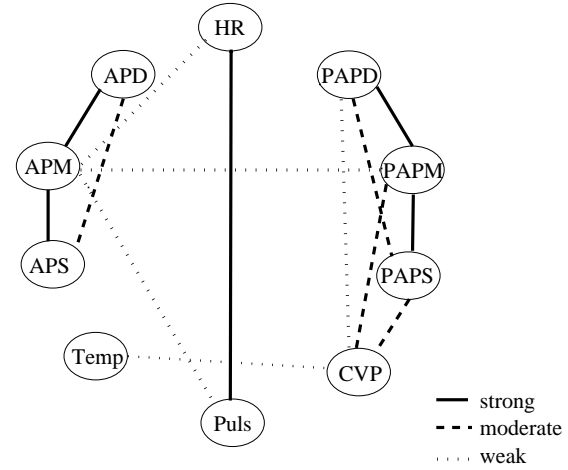
Figure 1: Partial correlation graph, one step selection. Different line types depict different strength of partial correlation.



The partial correlation graphs obtained from such a one-step selection generally match expected physiological relationships¹⁴. However, strong relationships may overlay weaker relationships so that the latter may be difficult to detect when estimating all partial linear relations jointly. Such masking effects can be overcome by applying a stepwise search strategy for model selection using separation properties of graphical models¹⁹. Starting from the initial classification obtained from the one-step selection, we checked this classification estimating the partial spectral coherencies in suitably chosen subgroups of the variables not changing the initial classification by more than one category¹⁹. Figure 2 shows the partial correlation graph resulting from the application of this stepwise search strategy for the same patient as before.

For all patients strong relationships could be identified between the arterial pressures (APS, APD, APM), between the pulmonary artery pressures (PAPS, PAPD, PAMP) and between heart rate and pulse, the strength of the relation between the systolic and the diastolic pressure being always lesser than between each of these and the corresponding mean pressure. Application of the stepwise search strategy revealed that CVP is strongest related to the pulmo-

Figure 2: Partial correlation graph, final selection. Different line types depict different strength of partial correlation.



nary artery pressures, while the temperature does not have strong relationships to any of the other variables. Hence, we can identify groups of strongly related variables from the partial correlation graphs. Such a partitioning of the variables into strongly related subgroups can be used for variable selection. The absence of edges between two groups of variables means that the variables in one of these groups do not add any information on the variables in the other group given the measurements of the remaining variables. On the other hand, if a variable has strong relationships to several other variables it provides a lot of information. Selecting APM from the strongly related subgroup of arterial pressures and neglecting APD and APSYS for clinical monitoring is therefore meaningful from a statistical point of view. The same applies to pulmonary artery pressures. Hence, we might end up with the selection PAMP, APM, HR and Temp based on the information obtained from the partial correlation graphs.

An alternative approach for dimension reduction is to extract latent variables from the observed time series which capture as much of the total variability as measured by the trace of the covariance matrix as possible. We scale the time series to unit variance and perform a standard principal component analysis based on correlations in order to explain as much of the total variability measured by the trace of the correlation matrix as possible.

Table 1 gives minimum (Min), maximum (Max), median (Med) and the upper and lower quartile (UQ and LQ) of the percentage of total variability explained by 1, 2, ..., 7 principal components for the 25 patients. We deduce that four principal components capture more than 90% of the total variability for half of the patients, and at least 85% for all of them, see the column Med and Min, respectively. Based on this crite-

tion we would probably use four principal components for an online monitoring, which is the minimal number of variables as suggested by the partial correlation graphs. After application of the automatic varimax rotation, which orthogonally transforms the directions of the extracted components to have many entries close to zero while spanning the same k-dimensional subspace, we can even associate the extracted latent variables with the groups found in the partial correlation graphs, see Table 2. However, these latent variables are still mixtures of all observed variables as all loadings are distinct from zero.

Table 1: Percentage of total variability explained by the first k components, $k = 1, 2, \dots, 7$.

k	Min	LQ	Med	UQ	Max
1	36.0	40.7	43.2	44.6	49.7
2	61.2	64.4	66.9	70.9	78.5
3	75.9	79.5	81.7	84.9	91.1
4	85.0	87.5	90.3	91.3	94.3
5	91.7	92.8	94.5	95.7	96.9
6	95.8	96.4	97.2	97.6	98.4
7	97.8	98.5	98.6	98.9	99.4

Table 2: Loading matrix after varimax rotation for one patient. The components can be identified with HR and PULS, the intrathoracic pressures, the arterial pressures and the temperature, respectively.

PAPS	-0.256	0.441	0.045	-0.275
PAPM	-0.071	0.533	0.019	-0.065
PAPD	0.074	0.493	0.064	0.180
CVP	0.172	0.508	-0.090	0.033
APS	0.123	0.112	0.419	0.056
APM	-0.032	-0.037	0.648	-0.016
APD	-0.056	-0.064	0.624	-0.034
HR	0.659	0.008	0.022	-0.023
PULS	0.659	0.010	0.022	-0.022
TEMP	-0.078	0.045	0.015	0.938

In order to further improve the interpretation of the extracted components we can extract one component from each group applying the simple dynamic factor model mentioned above separately to each group. Only from the group consisting of heart rate and pulse instead of extracting a latent variable we select the heart rate as its measurements, derived from ECG, are often more reliable.

In the following we compare the percentage of variability explained by a variable selection (VS) consisting of PAPM, APM, HR and Temp, by standard principal component analysis (PCA) and by the restricted

factor analysis (RFA). For this we regress the observed variables on the selected variables and on the extracted components, respectively. Then we investigate the total residual variance as well as the residual variance for all variables separately. For a selected variable this residual variance is zero, of course.

Table 3 shows the minimum, the maximum, the median and the lower and the upper quartile of the percentage of total variability explained by PCA, a variable selection and a restricted factor analysis for the same 25 patients as before. It can be deduced that the variable selection explains the main part of the total variability, at least about 80% in all cases considered here, see Min, but less than an ordinary PCA with the same number of components, of course. Performing a restricted factor analysis allows to regain some of this loss still providing meaningful latent variables. While the variable selection explains more than 85% of the total variability for half of the patients (Med), the extracted factors do so for about 75% of the patients (LQ), and the worst case (Min) almost increases to the lower quartile for the variable selection.

Table 3: Percentage of total variability explained by PCA, variable selection and restricted factor analysis.

	Min	LQ	Med	UQ	Max
PCA	85.0	87.5	90.3	91.3	94.3
VS	79.6	83.3	85.5	87.3	91.3
RFA	82.5	85.5	87.9	89.9	93.4

Table 4 shows that extracting one component from each group increases the explained variability for the variables not captured well by the variable selection substantially, see CVP for instance. At the same time the factors describe the variables included in the selection very well. When performing a PCA, however, the percentage of captured variability is about 80% or even more for 75% of the patients and each of the variables, which is quite high (see the column LQ).

Table 4: Percentage of variability of the individual variables explained by PCA, variable selection and a restricted factor analysis.

	LQ	Med	UQ	LQ	Med	UQ
	PAPS			PAPM		
PCA	81.4	85.6	89.8	94.3	95.8	97.4
VS	66.3	81.0	85.6	100.0	100.0	100.0
RFA	71.9	81.1	84.1	93.4	94.6	96.8
	PAPD			CVP		
PCA	84.7	88.6	91.3	79.8	84.6	90.5
VS	75.5	81.8	86.7	38.5	54.5	64.6
RFA	79.6	85.0	90.1	58.9	71.2	78.5

	HR			PULS		
PCA	93.6	94.8	97.5	92.9	95.5	97.4
VS	100.0	100.0	100.0	89.7	92.9	97.0
RFA	100.0	100.0	100.0	89.6	92.9	97.0
	APS			APM		
PCA	83.1	87.2	90.3	95.0	95.9	97.0
VS	63.4	75.1	81.2	100.0	100.0	100.0
RFA	74.4	83.5	87.0	95.8	96.5	96.8
	APD			TEMP		
PCA	83.9	87.7	92.1	83.9	87.6	92.6
VS	76.0	81.7	87.5	100.0	100.0	100.0
RFA	80.6	86.0	90.4	100.0	100.0	100.0

Conclusion

Statistical methods for dimension reduction aim at condensing the information provided by a high-dimensional time series into a few essential variables. In this regard, partial correlation graphs are a suitable tool since they help to explore the relations among the observable variables. The insights gained by this method can result in an improved online monitoring of vital signs as they allow an advanced application of dimension reduction techniques. One possibility is to select suitable subsets of important variables from the graphs. Alternatively, we can deduce information on the partial correlation structure from the partial correlation graph to enhance latent variable techniques. Deducing restrictions on the loading matrix from a graphical model means a compromise between variable selection and standard principal component analysis as the percentage of explained variability will typically be higher than for a variable selection and we get meaningful latent variables. In our study the groups of closely related variables obtained from the data analysis agree with the groups anticipated from medical knowledge. Although in more complex situations the results may be less clear-cut than in the situation considered here, we expect to gain reliable insights when applying this methodology to time series describing other variables, for which we have less medical knowledge so far.

References

1. Morris G, Gardner R. Computer applications. In: Hall J, Schmidt, G, Wood, L editors. Principles of critical care. McGraw-Hill: New York 1992, pp. 500-514.

2. Kohn LT, Corrigan JM, Donaldson M (eds): To Err Is Human. Building a Safer Health System. Committee on Quality of Health Care in America. Institute of Medicine. National Academy Press, Washington, DC, 2000.
3. Committee on Quality of Health Care in America, Institute of Medicine: Crossing the Quality Chasm: A New Health System for the 21st Century. National Academy Press, Washington, DC, 2001.
4. Whittaker J. Graphical Models in Applied Multivariate Statistics. Wiley: Chichester, 1990.
5. Cox DR, Wermuth N. Multivariate Dependencies. Chapman & Hall: London, 1996.
6. Lauritzen SL. Graphical Models. Clarendon Press: Oxford, 1996.
7. Edwards D. Introduction to graphical modelling. Second edition. Springer: New York, 2000.
8. Davies PL, Fried R, Gather U. Robust signal extraction for on-line monitoring data. J Statistical Planning and Inference 2003; to appear.
9. Fried R. Robust filtering of time series with trends. Preprint, Department of Statistics, University of Dortmund, Germany.
10. Brillinger DR. Remarks Concerning Graphical Models For Time Series And Point Processes. Revista de Econometria 1996; 16:1-23.
11. Dahlhaus R. Graphical interaction models for multivariate time series. Metrika, 2000; 51: 157-172.
12. Gather U, Imhoff M, Fried R. Graphical models for multivariate time series from intensive care monitoring. Statistics in Medicine 2002; 21: 2685-2701.
13. Brillinger DR. Time Series. Data Analysis and Theory. San Francisco: Holden Day 1981.
14. Imhoff M, Fried R, Gather U,. Detecting relationships between physiological variables using graphical modeling. Proc AMIA Symp 2002:340-344.
15. Reinsel GC. Elements of multivariate time series analysis. Second edition. Springer: New York 1997.
16. Peña D, Box GEP. Identifying a simplifying structure in time series. J. American Statistical Association 1987; 82: 836-843.
17. Fried R, Didelez V. Latent variable analysis and partial correlation graphs for multivariate time series. Preprint, Department of Statistics, University of Dortmund, Germany.
18. Dahlhaus R, Eichler M. Spectrum. C program, can be downloaded from <http://www.statlab.uni-heidelberg.de/projects>
19. Fried R, Didelez V. Decomposability and selection of graphical models for multivariate time series. Biometrika 2003; to appear.

Supported, in part, by the DFG (SFB475)